A Survey of Complexity Results for Kostka Numbers and Kronecker and Littlewood-Richardson Coefficients

Jonathan Liu, Max Ovsiankin

December 13, 2019

1 Introduction, Complexity Theory

Recall in class we talked about the *Littlewood-Richardson* coefficients, which are defined as structure constants for multiplication in the basis of Schur functions s_λ:

$$s_\mu s_\nu = \sum_\lambda c^\lambda_{\mu,\nu} s_\lambda$$

The Kroknecker coefficients $k^\lambda_{\mu,\nu}$ are defined in a similar relation between Schur functions:

$$\frac{1}{\prod_{i,j,k}(1 - x_i y_j z_k)} = \sum_{\lambda,\mu,\nu} k^\lambda_{\mu,\nu} s_\lambda(x) s_\mu(y) s_\nu(z)$$

In 1999, Stanley gave as a problem [1] to find a 'combinatorial interpretation' of the Kronecker coefficients $k^\lambda_{\mu,\nu}$, inspired by the combinatorial interpretation we have for $c^\lambda_{\mu,\nu}$. What does it mean for a collection of coefficients to have a 'combinatorial interpretation'?

Complexity theorists have an interesting answer to this question. Let $\langle\mu\rangle$ denote the number of bits used to represent the partition μ. The simplest representation is as a vector $(\mu_1, \ldots, \mu_k)$, so that $\langle\mu\rangle = |\mu|$.

Definition 1.1. Let us summarize the input $\mathbf{x} = (\lambda, \mu, \nu)$, so we can write $k^\lambda_{\mu,\nu} = k(\mathbf{x})$. A collection of coefficients $k(\mathbf{x})$ *has a* $\#\mathsf{P}$*-formula* if there exists a polynomial $p\colon \mathbb{N} \to \mathbb{N}$ and polynomial-time algorithm $F(\mathbf{x}, y)$ such that

$$k(\mathbf{x}) = \#\{y \in \{0,1\}^{p(\langle \mathbf{x}\rangle)} \mid F(\mathbf{x}, y) = 1\}$$

In other words, $k(\mathbf{x})$ counts the number of length-$p(\langle\mathbf{x}\rangle)$ bitstrings y for which $F(\mathbf{x}, y) = 1$.

The set of collections of coefficients with $\#\mathsf{P}$-formulas is called $\#\mathsf{P}$, which is an important complexity class in computational complexity theory. If a collection of coefficients has a $\#\mathsf{P}$-formula, a complexity theorist says that qualifies as a combinatorial interpretation, because the coefficients are counting the size of some class of objects, where we can verify if each individual object is in the class in polynomial-time.

coefficients	complexity of positivity	complexity of computation
$K_{\lambda,\mu}$ (Kostka numbers)	P	$\#\mathsf{P}$-complete†
$c^{\lambda}_{\mu,\nu}$ (Littlewood-Richardson coefficients)	P	$\#\mathsf{P}$-complete†
$k^{\lambda}_{\mu,\nu}$ (Kronecker coefficients)	NP-hard	$?^\dagger$
$\hat{k}^{\lambda}_{\mu,\nu}$ (restricted Kronecker coefficients)	NP-hard	$\#\mathsf{P}^\dagger$

Table 1: we will discuss in this talk any results marked with †

We can connect the complexity of computing the value of the coefficient with the complexity of a related problem:

Definition 1.2. The *positivity problem* for a collection of coefficients $k(\mathbf{x})$ is simply the question "Is $k(\mathbf{x}) > 0$?"

Consider a collection of of coefficients $k(\mathbf{x})$ that has a $\#\mathsf{P}$-formula. Then the positivity problem for that collection of coefficients is the same as:

$$k(\mathbf{x}) > 0 \iff \exists y \in \{0,1\}^{p(\langle \mathbf{x} \rangle)} \,\text{s.t.}\, F(\mathbf{x}, y) = 1$$

You may be familiar with the above definition: this means that deciding the positivity of $k(\mathbf{x})$ is in NP! In other words, deciding positivity of $k(\mathbf{x})$ is just answering the question 'does there exist a length-$p(\langle \mathbf{x} \rangle)$ bitstring for which $F(\mathbf{x}, y) = 1$?'. So we conclude

Lemma 1.3. If computing $k(\mathbf{x}) \in \#\mathsf{P}$, then the positivity problem $\{k(\mathbf{x}) > 0\} \in \mathsf{NP}$.

There are many, many more interesting NP problems, but the ones we will focus on will take this form. A summary of some results we will discuss is given in Table 1.

We haven't defined P and NP-complete, but many computer scientists believe that $\mathsf{P} \neq \mathsf{NP}$, which would mean that NP-complete problems are strictly harder than P problems. If $\mathsf{P} \neq \mathsf{NP}$, this would show that the complexity of deciding positivity of $k(\mathbf{x})$ is strictly harder than the complexity of deciding positivity of $c(\mathbf{x})$.

The Kostka numbers $K_{\lambda,\mu}$ count the number of tableaux of shape λ and content μ. We have already seen the Littlewood-Richardson and Kronecker coefficients $c^{\lambda}_{\mu,\nu}$, $k^{\lambda}_{\mu,\nu}$. Note that as stated in Table 1, a combinatorial interpretation for $k^{\lambda}_{\mu,\nu}$ is still unknown, but instead we will discuss a combinatorial interpretation for a restricted class $\hat{k}^{\lambda}_{\mu,\nu}$ from [2]. Before the innovations in [2], combinatorial interpretations were only known for collections of coefficients where the complexity of positivity is P. The combinatorial interpretation for $\hat{k}^{\lambda}_{\mu,\nu}$ was the first found for a collection of coefficients whose positivity was NP-hard.

Not only do we discuss that the complexity of computation for the Kostka numbers, Littlewood-Richardson coefficients, and restricted Kronecker coefficients is $\#\mathsf{P}$, we further show the former two are $\#\mathsf{P}$-complete, which means they are *the hardest possible counting problems.* In a very formal sense, a collection of coefficients $k(\mathbf{x})$ being $\#\mathsf{P}$-complete means that *any other* counting problem can be seen as a special case of $k(\mathbf{x})$ (you just need to find the right $\mathbf{x}$).

For a more in-depth discussion of the classes P, NP, $\#\mathsf{P}$, and completeness for NP and $\#\mathsf{P}$, see [3].

2 Littlewood-Richardson Coefficients

It turns out (as we discussed in class) that the Littlewood-Richardson Coefficients have a beautiful combinatorial interpretation called the *Littlewood-Richardson Rule* (see also [4], Theorem A.1.3).

Theorem 2.1 (Littlewood-Richardson Rule)**.** The Littlewood-Richardson coefficient $c^{\lambda}_{\mu,\nu}$ is equal to the number of *LR (Littlewood-Richardson)* semistandard Young tableaux of shape λ/μ, and type ν. A tableaux is LR if when its entries are read right to left, top to bottom, the number of copies of i that have been seen so far is at least the number of copies of $i+1$ that have been seen, for all i.

This interpretation is equivalent to the following that we will use for the rest of the presentation (see [5], page 61):

Corollary 2.2. The Littlewood-Richardson coefficient $c^{\nu}_{\lambda,\alpha}$ counts the number of LR tableau of shape $\lambda * \alpha$ of content ν.

Here, $\lambda * \alpha$ refers to the following operation producing a skew shape:

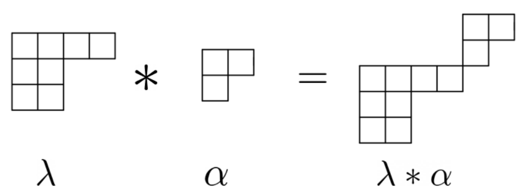

So $c^{\nu}_{\lambda,\alpha}$ still count LR tableau of skew shape, just a slightly different skew shape!

3 Hardness Results for Kotska and Littlewood-Richardson

We will cover the hardness results developed in [6].

Theorem 3.1. The computation of $K_{\lambda\mu}$ and $c^{\nu}_{\lambda\alpha}$ are #P-complete problems.

Proof Outline. We'll first demonstrate that computing LR coefficients is a #P-hard problem. Then, we demonstrate that Kostka numbers are a special case of the LR coefficients. Finally, we show that computing Kostka numbers is a #P-complete problem. Because any Kostka number can be computed by computing a LR coefficient, this means that computing LR coefficients is a #P-complete problem as well. □

Lemma 3.2. The problem of computing Littlewood-Richardson coefficients is #P-hard.

Proof. We note by Definition 1.1 that any problem #P is counting the size of some class of objects. Thus, because Corollary 2.2 gives us a combinatorial interpretation of the Littlewood-Richardson coefficients, it follows that the computation of these coefficients is in #P. □

One interesting note to make here is that we need to ensure that checking membership of the group takes polynomial time, as per Definition 1.1. Checking membership of the group is certainly linear in the sum of the entries because we only need to compute the addition of each entry in two different sums, but this turns out to potentially still be too slow; if we represent a tableau by its entries in binary, we can use as little as log(sum of entries), which means that our checking algorithm is suddenly exponential in the sum of the entries. Thus, we do need to refine our tableau to be unary, as in this case every tableau's size is exactly the sum of its entries.

Lemma 3.3. The computation of Kostka numbers is a #P-complete problem.

The definition of #P-complete allows us to readily use #P-complete problems to prove the completeness of other problems: if a #P-complete problem can be shown to be a special case of a different problem in #P, then that problem is by definition #P-complete as well! Before we can use this method, we need to introduce a #P-complete problem to use.

Definition 3.4 (Contingency Table)**.** For vectors $\mathbf{a} = (a_1, \ldots, a_m)$ and $\mathbf{b} = (b_1, \ldots, b_n)$, a contingency table is an $m \times n$ array with nonnegative integer entries where the sum of the entries in row i is a_i and the sum of the entries in column j is b_j.

Lemma 3.5. Denote $\mathbb{I}(\mathbf{a}, \mathbf{b})$ as the set of contingency tables with row sums $\mathbf{a}$ and column sums $\mathbf{b}$, with $\mathbf{a} \in \mathbb{Z}^2_{\geq 0}$ and $\mathbf{b} \in \mathbb{Z}^k_{\geq 0}$. Then the problem of computing $|\mathbb{I}(\mathbf{a}, \mathbf{b})|$, which is the number of $2 \times k$ contingency tables for given row and column sums, is #P-complete.

For a proof of Lemma 3.5, see [7].

Proof of Lemma 3.3. To demonstrate that computing Kostka numbers is #P-complete, we will demonstrate that the problem of counting $2 \times k$ contingency tables, which is #P-complete as mentioned earlier, can be reduced to computing Kostka numbers.

We are given $\mathbb{I}(\mathbf{a}, \mathbf{b})$, the set of $2 \times k$ contingency tables with row sums $\mathbf{a}$ and column sums $\mathbf{b}$. Set $\lambda = (a_1 + a_2, a_2)$ and $\mu = (b_1, \ldots, b_k, a_2)$. We claim that for this λ and μ,

$$|\mathbb{I}(\mathbf{a}, \mathbf{b})| = K_{\lambda\mu}.$$

To prove this, we'll construct a bijection between members of $|\mathbb{I}(\mathbf{a}, \mathbf{b})|$ and members of $K_{\lambda\mu}$.

First, using the RSK correspondence, we can construct a bijection between tables in $|\mathbb{I}(\mathbf{a}, \mathbf{b})|$ and pairs of SSYT (T_1, T_2) with a common shape ν, given by the set $\cup_\nu \mathbb{T}(\nu, \mathbf{a}) \times \mathbb{T}(\nu, \mathbf{b})$. In particular, for a $2 \times k$ contingency table M, each entry in T_1 comes from an entry in that row of M, so the contents of T_1 are actually given by $\mathbf{a}$. For the same reason, the contents of T_2 are given by $\mathbf{b}$.

Now, let's look at the possibilities for T_1. Because the content of T_1 is given by (a_1, a_2), we can reach the following conclusions about ν:

1. *ν contains at most two rows.* Because entries are strictly decreasing downwards, the fact that T_1 is filled solely with 1's and 2's means that it can not contain more than two rows.
2. *If the shape of ν is (ν_1, ν_2), then $\nu \trianglerighteq \mathbf{a}$.* Note that every 1 entry of an SSYT must be in the first row. Thus, we must have $\nu_1 \geq a_1$, so $\nu \trianglerighteq \mathbf{a}$.
3. *For any given ν and $\mathbf{a}$ such that $\nu \trianglerighteq \mathbf{a}$, there is exactly one possible SSYT with shape ν and content $\mathbf{a}$.* Note that the 1's must occupy the leftmost segment of the top row. Thus, there is only one possible configuration for the 1's so the 2's naturally fill in the remaining spaces.

Thus, even though we are given $(T_1, T_2) \in \cup_\nu \mathbb{T}(\nu, \mathbf{a}) \times \mathbb{T}(\nu, \mathbf{b})$, $\mathbb{T}(\nu, \mathbf{a})$ has size 1 when $\nu \trianglerighteq \mathbf{a}$ and 0 otherwise, so this set is in bijection with the set $\cup_{\nu \trianglerighteq \mathbf{a}} \mathbb{T}(\nu, \mathbf{b})$. Now, this is a union of tableaux with similar shapes and identical content, but Kostka numbers count the number of tableaux with a single given shape and content, so we must find a bijection from this union to a set of tableaux with identical shape and content.

To do so, we note that $\nu_1 \leq a_1 + a_2$ and $\nu_2 \leq a_2$. In a sense, we can view each ν as selecting $a_1 + a_2$ cells as a left-leaning subset of the shape $(a_1 + a_2, a_2)$, and filling it with $a_1 + a_2$ of the numbers 1 through k. Thus, a reasonable way to choose a new set of tableaux to encapsulate all viable tableaux is to take any viable tableau in $\cup_{\nu \trianglerighteq \mathbf{a}} \mathbb{T}(\nu, \mathbf{b})$, extend it to a tableau of shape $(a_1 + a_2, a_2)$, and fill in the added boxes with the value $k+1$. It's clear that this extension works for all possible ν, always adds a_2 values of $k+1$, always creates a tableau of size $(a_1 + a_2, a_2)$, and preserves the uniqueness of tableaux, so if we denote $\lambda = (a_1 + a_2, a_2)$ and $\mu = (b_1, \ldots, b_k, a_2)$, this process creates a bijection between $\cup_{\nu \trianglerighteq \mathbf{a}} \mathbb{T}(\nu, \mathbf{b})$ and $K_{\lambda\mu}$. □

We're not done yet! We can take this chain of bijections one step further, creating bijections between this form of Kostka numbers and a set of Littlewood-Richardson coefficients. In doing so, we demonstrate that computing $\mathbb{I}(\mathbf{a}, \mathbf{b})$ can also be reduced to computing a Littlewood-Richardson coefficient, demonstrating that it is also a #P-complete problem.

Lemma 3.6. The Kostka numbers created in the previous lemma, specifically $K_{\lambda\mu}$ for $\lambda \in \mathbb{Z}^2_{\geq 0}$ and $\mu \in \mathbb{Z}^{k+1}_{\geq 0}$, can be expressed as Littlewood-Richardson coefficients. In particular, there exist partitions α and ν such that $K_{\lambda\mu} = c^\nu_{\lambda\alpha}$.

Proof. Remember from Corollary 2.2 that the Littlewood-Richardson coefficient $c^\nu_{\lambda\alpha}$ is equal to the number of skew tableaux with shape $\lambda * \alpha$ and content ν. Thus, our task is to find a bijection between any tableau in $\mathbb{T}(\lambda, \mu)$ and some LR skew tableau with shape $\lambda * \alpha$ and content ν. The only real difference between the two sets is that the LR skew tableau need to have a skew shape, and need to follow the Littlewood-Richardson criteria, which means that reading from right to left, no number occurs more often than the number below it. In this case, a natural way to create a bijection is to simply start with the tableau in $\mathbb{T}(\lambda, \mu)$ and extend it by adding some tableau α that not only makes the tableau into a skew tableau but also ensures its LR-ness.

To this end, we set $\alpha = (\alpha_1, \ldots, \alpha_k)$ to have values

$$\alpha_i = \sum_{j>i} \mu_j,$$

and fill row i of α with entirely i. To adjust the content to match this, we set ν to have $\nu_i = \mu_i + \alpha_i$.

With this construction of α, we must verify the following properties:

- *The newly constructed skew tableau is LR.* Take any value i. Then when we read from right to left, we first read through α before reading through λ. In α, α_i occurs $\sum_{j>i} \mu_j$ times and then α_{i+1} occurs $\sum_{j>i+1} \mu_j$ times, so at the end of α we have that i has occurred μ_{i+1} times more than $i+1$. Therefore, because λ contains exactly μ_{i+1} instances of $i+1$, there can be no point where the count of $i+1$ is higher than the count of i. Every value i meets the criteria for LR, so this tableau is LR.

- *This construction is bijective.* If two tableaux λ are different, then clearly because we have only extended the tableaux, their extensions will be different as well. Furthermore, α is entirely decided by μ, so every extension will have the same shape and content, despite the tableaux all being different.

We have therefore created a bijection between each table in $\mathbb{T}(\lambda, \mu)$ and a corresponding LR skew tableau of shape $\lambda * \alpha$ and content ν. Thus, for this λ, α, ν, we have $K_{\lambda,\mu} = c^{\nu}_{\lambda,\alpha}$, as desired.

□

This completes our proof of Theorem 3.1. Because the #P-complete problem of counting $2 \times k$ contingency tables can be reduced through the bijections given above to the computation of LR coefficients, and the computation of LR coefficients is a problem in #P, it must be #P-complete as well. It's interesting to note that because computing Kostka coefficients is #P-complete, the problem of computing LR coefficients can also be reduced to computing Kostka coefficients.

4 Kronecker Coefficients

Now we detail some of the hardness results covered in [2]. It is still open to find a #P-formula for the Kronecker coefficients $k^{\lambda}_{\mu,\nu}$. So instead, we find a #P-formula for a restricted collection $\hat{k}^{\lambda}_{\mu,\nu}$ whose positivity is still NP-hard (we do not show the positivity is NP-hard here).

The restricted collection of coefficients we will use is triples (λ, μ, ν) that are *simplex-like*. The definition will actually involve geometric points sitting in a simplex, so we will need to do a couple of geometric things!

Let P_r be the size-r 3-dimensional simplex, i.e. $P_r = \{(x, y, z) \in \{0, \dots, r-1\}^3 \mid x+y+z \leq r-1\}$. Let $\mathrm{bary}(S) = \sum_{p \in S} p$ be the *barycenter* of the set S. Now consider a set P that lies 'in between' P_r and P_{r+1}, i.e. a set P where $P_r \subseteq P \subsetneq P_{r+1}$ such that $|P| = n$.

Lemma 4.1. For all P with $|P| = n$, such that there exists r with $P_r \subseteq P \subsetneq P_{r+1}$, we have

$$p(n) \triangleq (1,1,1) \cdot \mathrm{bary}(P) = (1,1,1) \cdot \mathrm{bary}(P_r) + r(n - |P_r|)$$

Proof. The projection of the barycenter of P to the diagonal $(1,1,1)$ can be computed as follows:

$$\begin{aligned}(1,1,1) \cdot \mathrm{bary}(P) = (1,1,1) \cdot \sum_{p \in P} p &= (1,1,1) \cdot \underbrace{\sum_{p \in P_r} p}_{x+y+z \leq r-1} + (1,1,1) \cdot \underbrace{\sum_{p \in P \setminus P_r} p}_{x+y+z=r} \\ &= (1,1,1) \cdot \mathrm{bary}(P_r) + r(n - |P_r|)\end{aligned}$$

This value only depends on n, so call the projection $p(n) = (1,1,1) \cdot \mathrm{bary}(P_r) + r(n - |P_r|)$. For any n, there is only one r such that there exist $|P| = n$ with $P_r \subseteq P \subsetneq P_{r+1}$, so we can define this value for all $n \geq 1$. □

Now consider an arbitrary point set in the cube $P \subseteq \{0, \dots, r-1\}^3$. It turns out that sets of the form we just considered minimize the projection of the barycenter:

Lemma 4.2. If $|P| = n$, $(1,1,1) \cdot \text{bary}(P) \geq p(n)$, with equality iff there exists r with $P_r \subseteq P \subsetneq P_{r+1}$.

Proof. Let r be the unique integer such that there exist $|P^*| = n$ with $P_r \subseteq P^* \subsetneq P_{r+1}$. Pick P^* so that $P \cap (P_{r+1} \setminus P_r) \subseteq P^* \cap (P_{r+1} \setminus P_r)$ (P^* contains all the points of P that are on the face of the size-$(r+1)$ simplex).

Consider all points $(x, y, z) \in P \setminus P^*$. If this set is empty, then $P = P^*$ so that $(1,1,1) \cdot \text{bary}(P) = p(n)$. Otherwise, note that any such (x, y, z) must necessarily have $x + y + z > r$. If we replace each (x, y, z) with a point in P^*, then $(1,1,1) \cdot \text{bary}(P)$ can only decrease (as $x + y + z \leq r$ for all points in P_{r+1}). So if this set is nonempty, we have that $(1,1,1) \cdot \text{bary}(P) > p(n)$. □

Now the partitions come in when counting types of these point sets. We say a point set P has *marginals* $(\lambda^T, \mu^T, \nu^T)$ if for all $0 \leq i \leq r-1$:

$$\lambda_i^T = \#\{(x,y,z) \in P \mid x = i\}, \mu_i^T = \#\{(x,y,z) \in P \mid y = i\}, \nu_i^T = \#\{(x,y,z) \in P \mid z = i\}$$

so that λ gives the marginal distribution of x-coordinates, μ for y-coordinates, and ν for z-coordinates. We call a point set P a *pyramid* if for any $(x, y, z) \in P$ and $0 \leq x' \leq x, 0 \leq y' \leq y, 0 \leq z' \leq z$, we have that $(x', y', z') \in P$ also.

Let $t^\lambda_{\mu,\nu}$ count the number of point sets P with marginals $(\lambda^T, \mu^T, \nu^T)$, and $p^\lambda_{\mu,\nu}$ count the number of pyramids with marginals $(\lambda^T, \mu^T, \nu^T)$. We have the following relation between $p^\lambda_{\mu,\nu}$, $k^\lambda_{\mu,\nu}$, and $t^\lambda_{\mu,\nu}$.

Lemma 4.3. For all (λ, μ, ν), we have $p^\lambda_{\mu,\nu} \leq k^\lambda_{\mu,\nu} \leq t^\lambda_{\mu,\nu}$.

The proof is very interesting, but involves far more representation theory than we have time to get into in this presentation. At least, it should be clear why $p^\lambda_{\mu,\nu} \leq t^\lambda_{\mu,\nu}$ (because all pyramids are trivially point sets).

Now we can define what it means to be simplex-like. Let λ^T be the transpose of the partition λ, then:

Definition 4.4. Let (λ, μ, ν) be such that $|\lambda| = |\mu| = |\nu| = n \neq 0$. Then (λ, μ, ν) is *simplex-like* if there exists some r such that the Young diagrams of λ, μ, ν have at most r+1 columns and

$$\sum_{i=0}^{r} i\lambda_i^T + \sum_{i=0}^{r} i\mu_i^T + \sum_{i=0}^{r} i\nu_i^T = p(n)$$

In the case of simplex-like marginals, it turns out that $p^\lambda_{\mu,\nu}$ and $t^\lambda_{\mu,\nu}$ coincide:

Lemma 4.5. If (λ, μ, ν) are simplex-like, then $p^\lambda_{\mu,\nu} = t^\lambda_{\mu,\nu}$.

Proof. Consider an arbitrary point set P with marginals $(\lambda^T, \mu^T, \nu^T)$. If we show that this point set is necessarily a pyramid, this establishes that $p^\lambda_{\mu,\nu} = t^\lambda_{\mu,\nu}$.

Let's look at the projection of the barycenter of P:

$$(1,1,1)\cdot \mathrm{bary}(P) = \sum_{p\in P}(1,1,1)\cdot p = \sum_{(x,y,z)\in P} x + \sum_{(x,y,z)\in P} y + \sum_{(x,y,z)\in P} z$$
$$= \sum_{i=0}^{r} i\lambda_i^T + \sum_{i=0}^{r} i\mu_i^T + \sum_{i=0}^{r} i\nu_i^T = p(n)$$

where the last equality follows because $(\lambda^T, \mu^T, \nu^T)$ are simplex-like. By Lemma 4.2, this means that there must exist r such that $P_r \subseteq P \subsetneq P_{r+1}$. But such a P is a pyramid! $\square$

Now let $\hat{k}^{\lambda}_{\mu,\nu}$ be the restricted Kronecker coefficients, which are the Kronecker coefficients, except nonzero iff (λ, μ, ν) are simplex-like. We can finally establish the following:

Theorem 4.6. $\hat{k}^{\lambda}_{\mu,\nu}$ has a #P-formula.

Proof. By Lemma 4.3 and Lemma 4.5, we have that $p^{\lambda}_{\mu,\nu} = k^{\lambda}_{\mu,\nu} = t^{\lambda}_{\mu,\nu}$ if (λ, μ, ν) are simplex-like. So we must only produce a #P-formula for $t^{\lambda}_{\mu,\nu}$, i.e. algorithm F such that

$$t^{\lambda}_{\mu,\nu} = \#\{y \in \{0,1\}^{p(\langle \mathbf{x}\rangle)} \mid F(\mathbf{x}, y) = 1\}$$

Like many algorithms complexity theorists produce, we will do this informally. Because (λ, μ, ν) are simplex-like and all have at most $r+1$ columns, we have $r \leq n+1$. It takes $3\log r$ bits to represent a point in $\{0, \ldots, r-1\}^3$, so a size-n point set can be represented by a bitstring of size $3n \log r \leq 3n^2$. Given a size-n point set P, the algorithm F will simply check that the marginals for P are $(\lambda^T, \mu^T, \nu^T)$, which just requires tallying up the coordinates for all the points, which can be done in polynomial time. $\square$

References

[1] Richard P. Stanley. Positivity problems and conjectures in algebraic combinatorics. In *in Mathematics: Frontiers and Perspectives*, pages 295–319. American Mathematical Society, 1999.

[2] Christian Ikenmeyer, Ketan Mulmuley, and Michael Walter. On vanishing of kronecker coefficients. *computational complexity*, 26:949–992, 2017.

[3] Sanjeev Arora and Boaz Barak. *Computational Complexity: A Modern Approach.* Cambridge University Press, New York, NY, USA, 1st edition, 2009.

[4] Richard P. Stanley and Sergey Fomin. *Symmetric Functions*, volume 2 of *Cambridge Studies in Advanced Mathematics*, page 286–560. Cambridge University Press, 1999.

[5] William Fulton. *Young Tableaux: With Applications to Representation Theory and Geometry.* London Mathematical Society Student Texts. Cambridge University Press, 1996.

[6] Hariharan Narayanan. On the complexity of computing kostka numbers and littlewood-richardson coefficients. *Journal of Algebraic Combinatorics*, 24(3):347–354, Nov 2006.

[7] Martin Dyer, Ravi Kannan, and John Mount. Sampling contingency tables. *Random Structures and Algorithms*, 10:487–506, 1997.